



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/775,101	02/11/2004	Walter James Reinhard	C697 0006/GNM	1573

720 7590 10/29/2008
OYEN, WIGGS, GREEN & MUTALA LLP
480 - THE STATION
601 WEST CORDOVA STREET
VANCOUVER, BC V6B 1G1
CANADA

EXAMINER

SIDDIQI, MOHAMMAD A

ART UNIT	PAPER NUMBER
----------	--------------

2454

MAIL DATE	DELIVERY MODE
-----------	---------------

10/29/2008

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

DETAILED ACTION

1. Claims 1-46 are presented for examination.

Claim Rejections - 35 USC § 102

2. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in a patent granted on an application for patent by another filed in the United States before the invention thereof by the applicant for patent, or on an international application by another who has fulfilled the requirements of paragraphs (1), (2), and (4) of section 371(c) of this title before the invention thereof by the applicant for patent.

3. Claims 1-11 and 19-46 are rejected under 35 U.S.C. 102(e) as being anticipated by Pettey et al. (7,046,668) (hereinafter Pettey).

4. As per claim 1, Pettey discloses a method for exchanging data between compute nodes of a computer system comprising:

a plurality of compute nodes interconnected (PCI Express, fig 3, col 11, lines 4-26) by an inter-node communication network (fig 1, col 7, 41-47), each of the compute nodes having an independent address space (own memory structure, fig 1, col 7, lines 54-67) and comprising:

a local packetized interconnect (col 10, lines 38-43),

a network interface (col 8, line 10) coupled to the local packetized interconnect and the inter-node communication network (col 10, lines 38-43), at least one data processor coupled to the local packetized interconnect (col 10, lines 30-43; col 11, lines 4-26); and,

a memory system coupled to the local packetized interconnect (col 9, lines 55-65); the method comprising tunneling data from the sending compute node to the receiving compute node by (col 23, lines 35-40): placing a local packetized interconnect packet on the local packetized interconnect of the sending compute node (col 11, lines 4-26; col 15, lines 9-15);

receiving the local packetized interconnect packet at the network interface of the sending compute node (col 18, line 64 – col 19, line 11);

encapsulating the local packetized interconnect packet in an inter-node communication network packet addressed to the receiving compute node (col 18, line 64 – col 19, line 11; col 23, lines 35-40);

dispatching the inter-node communication network packet to the receiving compute node by way of the inter-node communication network (col 18, line 64 – col 19, line 11);

receiving the inter-node communication network packet at the network interface of the receiving compute node (col 18, line 64 – col 19, line 36);

extracting the local packetized interconnect packet from the inter-node communication network packet (col 18, line 64 – col 19, line 36); and,

placing the extracted packet onto the local packetized interconnect of the receiving compute node (PCI Express, col 18, line 64 – col 19, line 36)

wherein, when the local packetized interconnect packet is on the local packetized interconnect of the sending compute node (PCI Express, col 18, line 64 – col 19, line 36), portions of the local packetized interconnect packet other than any addresses and any check values (fig 11, col 18, line 64 – col 19, line 36) are the same as the corresponding portions of the local packetized interconnect packet when the local packetized interconnect packet is on the local packetized interconnect of the receiving compute node (PCI Express, col 18, line 64 – col 19, line 36).

5. As per claim 2, Pettey discloses associating a first range of addresses in an address space of a sending one of the compute nodes with the network interface of the sending compute node (col 16, lines 7-31) and, at the network interface of the sending compute node associating the first range of addresses with the receiving compute node wherein the method comprises determining that the local packetized interconnect packet is associated with an address in the first range of addresses upon receiving the local packetized interconnect packet at the network interface of the sending compute node (col 16, lines 7-31; col 18, line 64 – col 19, line 36).

6. As per claim 3, Pettey discloses the method comprises performing an address translation on the local packetized interconnect packet after receiving the packet at the network interface of the sending compute node and prior to placing the extracted packet

onto the local packetized interconnect of the receiving compute node (col 18, line 64 – col 19, line 36).

7. As per claim 4, Pettey discloses performing the address translation comprises editing the local packetized interconnect packet by changing an address to which the packet is addressed from an address in the first range of addresses to a corresponding address in an address space of the receiving compute node (PCI Express logic, col 16, lines 7-31).

8. As per claim 5, Pettey discloses the address translation is performed at the network interface of the sending compute node (col 16, lines 7-31).

9. As per claim 6, Pettey discloses allocating a region of the memory of the receiving compute node to receive data from the sending compute node, memory locations in the allocated region being addressable by addresses in a second range of addresses corresponding to the first range of addresses, and communicating the second range of addresses from the receiving compute node to the sending compute node prior to tunneling the data from the sending compute node to the receiving compute node (PCI configuration logic, col 16, lines 7-31; col 18, line 64 – col 19, line 36).

10. As per claim 7, Pettey discloses at the sending compute node, prior to tunneling the data, computing an address transformation between the first and second address ranges and using the address transformation to perform the address translation (col 16, lines 7-31; col 18, line 64 – col 19, line 36).

11. As per claim 8, Pettey discloses the address transformation comprises an address translation table (col 16, lines 7-31).

12. As per claim 9, Pettey discloses the first and second ranges of addresses are equal in size (col 16, lines 7-31).

13. As per claim 10, Pettey discloses the address translation is performed at the network interface of the receiving compute node (col 16, lines 7-31; col 18, line 64 – col 19, line 36).

14. As per claim 11, Pettey discloses the local packetized interconnect packet comprises a write request packet comprising data to be written to a memory location in the memory system of the receiving compute node corresponding to the address and the method comprises writing the data to the memory location in the receiving compute node (shared I/O, col 18, line 64 – col 19, line 36).

15. As per claim 19, Pettey discloses the local packetized interconnect packet comprises a read request packet (PCI Express logic, col 16, lines 7-31).

16. As per claim 20, Pettey discloses at the network interface of the sending compute node, associating the first range of addresses with a plurality of receiving compute nodes (col 16, lines 7-31).

17. As per claim 21, Pettey discloses maintaining a correspondence between the first range of addresses and address ranges in address spaces of the plurality of receiving compute nodes (col 16, lines 7-31).

18. As per claim 22, Pettey discloses the address ranges in the address spaces of the plurality of receiving compute nodes are not all the same (col 16, lines 7-31).

19. As per claim 23, Pettey discloses dispatching the inter-node communication network packet to the plurality of receiving compute nodes by way of a multicast feature of the inter-node communication network (Ethernet protocol, col 6, lines 5-24).

20. As per claim 24, Pettey discloses at the sending compute node, encapsulating each of a plurality of copies of the local packetized interconnect packet in an inter-node communication network packet addressed to a different one of the plurality of receiving compute nodes and dispatching each of the inter-node communication network packets

to the corresponding receiving compute node by way of the inter-node communication network (col 16, lines 7-31).

21. As per claim 25, Pettey discloses placing a plurality of local packetized interconnect packets addressed to one or more addresses in the first range of addresses on the local packetized interconnect of the sending compute node wherein encapsulating the packet in the inter-node communication network packet comprises encapsulating the plurality of local packetized interconnect packets in the same inter-node communication network packet (PCI Express logic, col 16, lines 7-31).

22. As per claim 26, Pettey discloses the plurality of local packetized interconnect packets each comprise header information and a payload, at least some of the header information is the same for each of the local packetized interconnect packets and encapsulating the plurality of local packetized interconnect packets comprises placing into the inter-node communication network packet the payloads of all of the plurality of local packetized interconnect packets and the header information from fewer than all of the plurality of local packetized interconnect packets (PCI Express logic, col 16, lines 7-31).

23. As per claim 27, Pettey discloses the sending and receiving compute nodes are peers (fig 1, switch, col 16, lines 7-31).

24. As per claim 28, Pettey discloses the sending compute node has substantially the same construction as the receiving compute node (col 16, lines 7-31).

25. As per claim 29, Pettey discloses the local packetized interconnect packet comprises an atomic read-modify-write packet (col 14, lines 32).

26. As per claim 30, Pettey discloses the local packetized interconnects of the compute nodes operate at data rates in excess of 300 MBps (PCI Express bus, col 11, lines 4-27).

27. As per claim 31, Pettey discloses the inter-node communication network is characterized by a link data rate of at least 1 GBps (PCI Express bus, col 11, lines 4-27).

28. As per claim 32, Pettey discloses at the receiving compute node, altering a correspondence of the first range of addresses to addresses in the address space of the receiving compute node by changing the second range of addresses to a third range of addresses (Domains, col 16, lines 7-31).

29. As per claim 33, Pettey discloses the packet placed on the local packetized interconnect of the sending compute node is a read response packet and the method comprises converting the read response packet into a write request packet (col 18, lines

55 – col 19, line 36).

30. As per claim 34, Pettey discloses issuing a read request directed to a location in the memory system of the sending compute node wherein, placing the packet on the local packetized interconnect of the sending compute node is performed by the memory system of the sending compute node in response to the read request (col 18, lines 55 – col 19, line 36).

31. As per claim 35, Pettey discloses allocating a memory ID to a range of addresses in an address space of the receiving compute node and associating a corresponding first range of addresses in an address space of the sending compute node with the memory ID and the receiving compute node; at the sending compute node, determining an offset from an address associated with the local packetized interconnect packet and the first range of addresses and passing the memory ID and the offset in the inter-node communication network packet to the receiving compute node (col 18, lines 55 – col 19, line 36).

32. As per claim 36, Pettey discloses at the network interface of the receiving compute node, determining from the offset and memory ID a destination address in the range of addresses in the address space of the receiving compute node wherein placing the extracted packet onto the local packetized interconnect of the receiving compute node comprises addressing the extracted packet to the destination address (col 18,

lines 55 – col 19, line 36).

33. As per claim 37, Pettey discloses establishing a reciprocal pathway for tunneling data packets from the receiving compute node to the sending compute node (col 18, lines 55 – col 19, line 36).

34. As per claim 38, Pettey discloses at the network interface of the receiving compute node, receiving a request from the data processor of the receiving compute node that data be obtained from the sending compute node and, in response thereto, generating a local packetized interconnect read request packet, encapsulating the read request packet, and forwarding the read request packet to the sending compute node by way of the inter-node communication network wherein the local packetized interconnect packet is a read response packet generated in response to the read request packet (PCI Express, col 18, lines 55 – col 19, line 36).

35. As per claims 39-42, 43-44, and 45-46, claims are rejected for the same reasons as claims 1-4, above.

Claim Rejections - 35 USC § 103

36. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are

such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

37. Claims 12-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Pettey et al. (7,046,668) (hereinafter Pettey) in view of Carnevale et al. (7,024,613) (hereinafter Carnevale).

38. As per claim 12, Pettey does not explicitly show how the write request packet comprises a write confirmation request and the method comprises, detecting the write confirmation request at the network interface of the sending compute node, generating a write confirmation packet at the network interface of the sending compute node and dispatching the write confirmation packet on the local packetized interconnect of the sending compute node. Carnevale, in the same field of endeavor, discloses the write request packet comprises a write confirmation request and the method comprises, detecting the write confirmation request at the network interface of the sending compute node, generating a write confirmation packet at the network interface of the sending compute node and dispatching the write confirmation packet on the local packetized interconnect of the sending compute node (col 4, lines 1-24). Therefore, it would have been obvious for one of ordinary skill in the art at the time the invention was made to modify Pettey in view of Carnevale by including a buffer management logic, because

Carnevale suggests to achieve a desired latency in the interconnect fabric.

One of ordinary skill in the art would have been motivated to modify Pettey in view of Carnevale by including buffer management logic to achieve a desired latency in the interconnect fabric.

39. As per claim 13, the claim is rejected for the same reasons as claim 12, above. In addition, Carnevale discloses the network interface of the sending compute node, retaining a copy of the write request packet, maintaining a write completion timer and, if the write completion timer times out (col 4, lines 1-24) using the copy of the write request packet to resend the write request packet to the receiving compute node by way of the inter-node communication network (retransmit, col 4, lines 1-24).

40. As per claim 14, the claim is rejected for the same reasons as claim 12, above. In addition, Carnevale discloses the write request packet comprises a write confirmation request and the method comprises: detecting the write confirmation request at the memory system of the receiving compute node, placing on the local packetized interconnect of the receiving compute node a write confirmation packet containing the address of the memory location in the address space of the receiving compute node; changing the write confirmation packet by altering the contained address of the memory location in the address space of the receiving compute node to a corresponding address in the first range of addresses; and, subsequently placing the write confirmation packet on the local packetized interconnect of the sending compute node (col 4, lines 1-

24).

41. As per claim 15, the claim is rejected for the same reasons as claim 12, above. In addition, Carnevale discloses changing the write confirmation packet is performed at the network interface of the sending compute node (col 4, lines 1-24).

42. As per claim 16, the claim is rejected for the same reasons as claim 12, above. In addition, Carnevale discloses changing the write confirmation packet is performed at the network interface of the receiving compute node (col 4, lines 1-24).

43. As per claim 17, the claim is rejected for the same reasons as claim 12, above. In addition, Carnevale discloses the write request packet comprises a write confirmation request and the method comprises, detecting the write confirmation request at the network interface of the receiving compute node, generating a write confirmation packet at the network interface of the receiving compute node and dispatching the write confirmation packet to the sending compute node (col 4, lines 1-24).

44. As per claim 18, the claim is rejected for the same reasons as claim 12, above. In addition, Carnevale discloses at the network interface of the receiving compute node, retaining a copy of the write request packet, maintaining a write completion timer and, if the write completion timer times out, using the copy of the write request packet to

resend the write request packet to the memory system of the receiving compute node
(col 4, lines 1-24).

Response to Arguments

45. Applicant's arguments filed 07/23/2008 have been fully considered but they are not persuasive, therefore rejections to claims 1-44 is maintained.

46. Examiner has cited particular columns and line numbers in the references as applied to the claims above for the convenience of the applicant. Although the specified citations are representative of the teachings of the art and are applied to the specific limitations within the individual claim, other passages and figures may apply as well. It is respectfully requested from the applicant in preparing responses, to fully consider the references in entirety as potentially teaching all or part of the claimed invention, as well as the context of the passage as taught by the prior art or disclosed by the Examiner. In general, Applicant's arguments reflect a difference of opinion over the teachings of the prior art and how these teachings would be evaluated in light of the knowledge generally available to those in the appropriate art and the level of ordinary skill in the art. Moreover, Applicant's take an overly narrow view of the claim language.

47. In the remarks applicants argued that:

Argument: Pettey does not a plurality of compute nodes interconnected by an inter-node communication network, each of the compute nodes having an independent address space and comprising: encapsulating the local packetized interconnect packet in an inter-node communication network packet addressed to the receiving compute node.

Response: Pettey disclose a plurality of compute nodes (102, 104, 108, fig 1) interconnected (122, 124, 126, fig 1 ; PCI Express Switch, 322, fig 3, col 11, lines 4-26) by an inter-node communication network (122, 124, 126, fig 1, col 7, 41-47), each of the compute nodes having an independent address space (own memory structure, fig 1, col 7, lines 54-67) and comprising: encapsulating the local packetized interconnect packet (PSI Express + packet for shared I/O; fig 9-11) in an inter-node communication network packet addressed (fig 11) to the receiving compute node (col 18, line64 – col 19, line 11; col 23, lines 35-40, “ In one embodiment, the contents of the OS Domain Header are first established by the shared I/O switch 610 by embedding the port number in the shared I/O switch 610 that is coupled to the upstream root complex from which a packet originated, or for which a packet is intended, as the OS Domain Number.”).

Argument: Pettey does not disclose associating a first range of addresses in an address space of a sending one of the compute nodes with the network interface of the sending compute and, at the network interface of the sending compute node associating the first range of addresses with the receiving compute node wherein the method comprises determining that the local packetized interconnect packet is associated with an address in the first range of addresses upon receiving the local packetized interconnect packet at the network interface of the sending compute node .

Response: Pettey discloses associating a first range of addresses in an address space of a sending one of the compute nodes with the network interface of the sending compute node (tables associated with domain, col 16, lines 7-31) and, at the network interface of the sending compute node associating the first range of addresses with the receiving compute node wherein the method comprises determining that the local packetized interconnect packet is associated with an address (Fig 11) in the first range of addresses upon receiving the local packetized interconnect packet at the network interface of the sending compute node (col 16, lines 7-31; col 18, line 64 – col 19, line 36, “the present invention provides an apparatus for associating packets in a load/store serial communication fabric with root complexes to allow the root complexes to share an input/output (I/O) endpoint. The apparatus includes a shared I/O switch and a link. The shared I/O switch is coupled to each of the root complexes and has

routing control to associate the packets from each of the root complexes with the root complex they originate from by incorporating a field within the packets. The link is between the shared I/O switch and the input/output (I/O) endpoint. The link allows the packets to be transferred from the shared I/O switch to the input/output (I/O) endpoint with the field. The input/output (I/O) endpoint associates the packets with their associated root complexes by examining the field..) .

Applicant's arguments on pages 20-29 amount to a general allegation that the claims define a patentable invention without specifically pointing out how the language of the claims patentably distinguishes them from the references. Applicant argues that the prior art of Pettey does not teach address translation; allocating a region of the memory to the computing node; the read/write request of packets; dispatching the inter-node communication network packet to the plurality of receiving compute nodes by way of a multicast feature of the inter-node communication network ; altering a correspondence of the first range of addresses to addresses in the address space of the receiving compute node by changing the second range of addresses to a third range of addresses. Examiner respectfully disagrees. Pettey teaches address translation (col 22, lines 44-54, "a shared I/O switch to associate and route packets from root complexes to their associated endpoints"); allocating a region of the memory to the computing node (col 23, lines 14-34, inherent in the operating system); the read/write request of packets (col 23. lines 14-34; Shared I/O); dispatching the inter-node

communication network packet to the plurality of receiving compute nodes by way of a multicast feature of the inter-node communication network (fig 2); altering a correspondence of the first range of addresses to addresses in the address space of the receiving compute node by changing the second range of addresses to a third range of addresses (Multi server computing node 202, fig 2, please see summary of the invention rack, within a serial load/store fabric, from a plurality of root complexes with their originating root complex, to allow the plurality of root complexes to share an I/O endpoint ; The shared I/O [read/write] switch [interconnect] is coupled to each of the root complexes and has routing control to associate the packets from each of the root complexes with the root complex they originate from by incorporating a field within the packets).

Conclusion

48. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after the end of the **THREE-MONTH** shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to MOHAMMAD A. SIDDIQI whose telephone number is (571)272-3976. The examiner can normally be reached on Monday -Thursday.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Nathan J. Flynn can be reached on (571) 272-1915. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Nathan J. Flynn/
Supervisory Patent Examiner, Art Unit 2454